DTLS Performance in Duty-Cycled Networks

Mališa Vučinić[‡*¶], Bernard Tourancheau[*], Thomas Watteyne[§],
Franck Rousseau[*], Andrzej Duda[*], Roberto Guizzetti[¶], and Laurent Damon[¶]
[‡]BSAC, University of California, Berkeley, CA, USA.
[*]Grenoble Alps University, CNRS Grenoble Informatics Laboratory UMR 5217, Grenoble, France.
[¶]STMicroelectronics, Crolles, France.
[§]Inria, EVA Team, Rocquencourt, France.
Email: firstname.lastname@{eecs.berkeley.edu, imag.fr, st.com, inria.fr}

arXiv:1507.05810v1 [cs.NI] 21 Jul 2015

***Abstract*—The Datagram Transport Layer Security (DTLS) protocol is the IETF standard for securing the Internet of Things. The Constrained Application Protocol, ZigBee IP, and Lightweight Machine-to-Machine (LWM2M) mandate its use for securing application traffic. There has been much debate in both the standardization and research communities on the applicability of DTLS to constrained environments. The main concerns are the communication overhead and latency of the DTLS handshake, and the memory footprint of a DTLS implementation. This paper provides a thorough performance evaluation of DTLS in different duty-cycled networks through real-world experimentation, emulation and analysis. In particular, we measure the duration of the DTLS handshake when using three duty cycling link-layer protocols: preamble-sampling, the IEEE 802.15.4 beacon-enabled mode and the IEEE 802.15.4e Time Slotted Channel Hopping mode. The reported results demonstrate surprisingly poor performance of DTLS in radio duty-cycled networks. Because a DTLS client and a server exchange more than 10 signaling packets, the DTLS handshake takes between a handful of seconds and several tens of seconds, with similar results for different duty cycling protocols. Moreover, because of their limited memory, typical constrained nodes can only maintain 3-5 simultaneous DTLS sessions, which highlights the need for using DTLS parsimoniously.**

I. Introduction

Secure Sockets Layer (SSL) and its successor Transport Layer Security (TLS) [1] are fundamental cryptographic protocols supporting secure communication over the Internet. With the advent of the Internet of Things (IoT) and its applications, we face the problem of supporting communication security for IoT devices that present stringent energy, memory, and CPU constraints. Datagram TLS (DTLS) [2] is a version of TLS running over UDP used by the Constrained Application Protocol (CoAP) [3], ZigBee IP and Lightweight Machine to Machine (LWM2M) (other standards for constrained IoT networks) to secure IoT application traffic.

Yet, running DTLS on constrained IoT devices is challenging, in particular because of the amount of traffic needed to establish a DTLS session and of the memory footprint of a DTLS implementation [4], [5]. DTLS benchmarks exist [6], [7] and focus on memory footprint and message size for different cipher suites.

To achieve long lifetimes up to several years on a battery, IoT devices use duty cycling link-layer protocols—they follow a sleep/wakeup schedule to minimize the time their radio transceivers are on, which reduces the energy consumption.

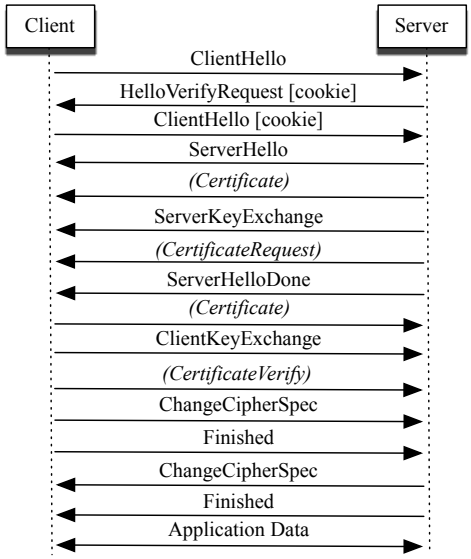

Fig. 1. Message exchange during a DTLS handshake. Messages in parentheses are not sent for pre-shared key cipher suites.

On the other hand, duty cycling results in higher latency and lower throughput, which has a direct impact on the DTLS performance. The goal of this paper is to provide a thorough evaluation of the DTLS performance on top of representative duty-cycled networks. More specifically, the contributions of this paper are the following:

- We measure the duration of the DTLS handshake and energy consumption for the following three duty cycling protocols: preamble sampling [8], IEEE 802.15.4 beacon-enabled mode [9], and IEEE 802.15.4e Time Slotted Channel Hopping (TSCH) [10]. We use several evaluation methods (emulation, measurements on a real-world testbed, an analysis) to obtain meaningful results. This part is the core of the paper.
- We quantify the impact of the limited memory on the number of simultaneous DTLS sessions a constrained IoT device can maintain.
- We show that the probability for a DTLS session establishment to fail because the server runs out of memory to hold the associated state can be modeled with the Engset loss formula.

The remainder of this paper is organized as follows. Section II gives an overview of the DTLS protocol. Sections III-A, III-B and III-C present the experimental and analytical performance results of DTLS on top of duty cycling protocols. Section IV discusses the impact of memory constraints on DTLS. Section V presents related work and Section VI concludes the paper.

II. Overview of DTLS

Securing application traffic is often achieved by transferring data over a secure channel between the two communicating end-points. In the network stack, this secure end-to-end channel can be established at the network layer (e.g. IPsec), the transport layer (e.g. TLS), or the application layer (e.g. SSH). For application development, security at the Transport layer is the most common. The *de-facto* security standard for the Internet application traffic is Secure Sockets Layer (SSL) and its IETF successor Transport Layer Security (TLS) [1]. TLS was designed for client-server applications that operate over a reliable transport. To establish a secure channel, a client and a server first perform the TLS handshake during which they authenticate each other and derive the symmetric keys to use during the session.

DTLS [2] is an extension of TLS for datagram transport and runs over UDP rather than TCP. Like TLS, DTLS protects the payload with encryption and authentication. DTLS records are 8 bytes longer than in TLS, as DTLS adds an explicit 8-byte sequence number. Stream ciphers, such as RC4, create an inter-record cryptographic context that introduces vulnerabilities with dropped and reordered packets. Consequently, DTLS bans the use of stream ciphers and relies on block ciphers for encrypting and authenticating records.

All messages carried by DTLS are encapsulated within DTLS *records* that add a constant 13 byte overhead per datagram. The Record Layer supports four DTLS upper protocols: 1) *Handshake* protocol establishing a secure authenticated session between two peers, negotiating algorithms, and the key material; 2) *Alert* protocol signaling session closure or eventual errors; 3) *Change Cipher Spec* protocol signaling modifications to encryption strategies; and 4) *Application Data* protocol carrying application data. To deal with Denial of Service (DoS) attacks, the Handshake protocol uses a stateless *cookie* exchange: before the server allocates any resources, the client needs to resends the *cookie* thus proving that the client can receive messages at a given IP address.

Fig. 1 shows a message exchange during a DTLS handshake. Once the client has replayed the stateless *cookie* from the server `HelloVerifyRequest` message, the server allocates the necessary resources. It chooses its preferred cipher from the client cipher set and notifies the client using a `ServerHello` message. The messages exchanged during key negotiation depend on the cipher. When using a pre-shared key, the message containing the server certificate is omitted and the server optionally sends `ServerKeyExchange` indicating to the client which pre-shared key to use. In this case, the two parties authenticate each other with the common secret (also used to derive the session keys).

Because of different application types, DTLS has been used differently in IoT networks and on the traditional Internet. It is very common for a regular Internet host to establish short-lived TLS sessions, for example when browsing the Internet `https` URLs. An IoT device typically periodically reports sensor readings to a server and therefore uses one long-lived DTLS session, which is a good thing as constrained IoT networks cannot handle frequent expensive DTLS handshakes, as highlighted in this paper. However, we are witnessing the emergence of applications in which mobile workers establish short-lived DTLS sessions with individual nodes using handheld devices, for example for maintenance or drive-by metering [11]. In this context, it is important to understand the limitations of DTLS when duty cycling protocols are used.

Duty cycling is a cornerstone technique for achieving long lifetimes of IoT devices. A typical IoT node with an IEEE 802.15.4 radio will deplete a 2200 mAh AA battery in about a week, if the radio is left on continuously (either receiving or transmitting). State-of-the-art duty cycling protocols reduce the radio duty cycle below 1%, thereby extending the device lifetime to several years. The price of such aggressive duty cycling is an increased network delay and reduced throughput. In the following sections, we study the effects of duty cycling protocols on the performance of DTLS through emulation, real-world experimentation, and analysis.

III. DTLS Performance in Duty-Cycled Networks

A. Preamble Sampling Protocols

To reduce radio idle listening, nodes in preamble sampling protocols periodically wake up to check the state of the wireless medium. The period at which a node wakes up is called the "Check Interval" (`CI`); the lower the `CI`, the more often nodes check the medium, and the higher their idle radio duty-cycle. Before sending a data packet, the transmitter repeatedly transmits a special control frame (called *strobe*) for at least the `CI`. If the receiver receives such a frame when checking the medium, it leaves its radio on until it receives the data frame. In this section, we use X-MAC [8], [12], arguably the most popular preamble sampling MAC protocol. X-MAC adds the receiver address in each *strobe*, so only the destination keeps its radio on to receive the data. All nodes (client, server and relay nodes) use the same `CI` value.

We leverage tinyDTLS[1], an open-source DTLS implementation and its port to the Contiki operating system [13]. We use a pre-shared key cipher suite of DTLS with AES operating in CCM mode with 8-byte long authentication tags (`TLS_PSK_WITH_AES_128_CCM_8`). We evaluate the performance of this implementation by emulation, using the instruction-level MSP430 emulator MSPSim, and the Contiki Cooja simulator [14]. We emulate WiSMote, a popular constrained IoT platform featuring a 16-bit MSP430 microcontroller with 16 kB of RAM running at 12 MHz, and the

[1] http://tinydtls.sourceforge.net/

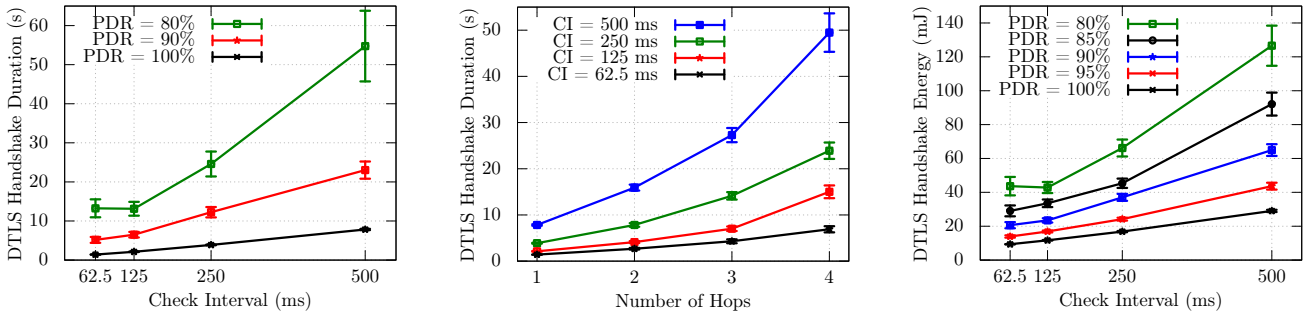

(a) DTLS handshake duration, single-hop.

(b) DTLS handshake duration, multi-hop (PDR=100%).

(c) Energy, single-hop.

Fig. 2. Cost of a DTLS handshake in preamble sampling protocols.

IEEE 802.15.4-compliant CC2520 radio. The binary file used for emulation runs on WiSMote hardware.

In the Cooja simulator, we create a linear network of 2 to 5 nodes, depending on the required number of hops. The DTLS client (on one end of the linear network) repeatedly performs the handshake with the DTLS server (on the other end). There is no other traffic in the network. We estimate the energy consumption using Energest, a Contiki per-component profiling tool that measures the consumption of both the microcontroller and radio. We average the results over 1000 DTLS handshakes and present with a 95% confidence interval.

Overall results. Fig. 2 shows the measured average handshake duration and the energy consumption when DTLS runs on preamble sampling protocols, in the single/multi-hop case, and for different values of `CI` and link PDR[2]. The energy consumption in Fig. 2(c) is that of the DTLS client (running at 2.8 V), during the DTLS handshake. Although absolute energy values are specific to WiSMote, this platform is representative of hardware commonly deployed today, and the trends in Fig. 2(c) apply to all platforms. Overall, a DTLS handshake takes 1–50 s, with an energy consumption in 10–200 mJ range.

Impact of `CI`. At PDR=100%, the DTLS handshake duration and energy consumption grows linearly with `CI`. This is expected, as a larger `CI` reduces the rate at which nodes can exchange packets (hence a longer duration).

Impact of the number of hops. Similarly, separating DTLS server and client by additional hops increases the duration of the handshake. For PDR=100%, the increase is linear (some retransmissions still occur due to the hidden terminal problem); for PDR<100%, the increase is faster as a packet can be lost on each of the hops.

Impact of PDR. In any wireless environment, external interference and multi-path fading cause the PDR to be below 100%. When a DTLS message is lost, a timeout event occurs at the DTLS layer, which triggers retransmission of DTLS messages. X-MAC implicitly recovers from lost *strobes*, but does not detect failed receptions of data frames as there are no link-layer acknowledgements. This means that, when a DTLS packet is lost, the DTLS implementation waits for 2 s (the default tinyDTLS timeout value) before resending, causing a longer latency. Dropping the PDR from 100% to 90% roughly triples the handshake duration (Fig. 2(a)) and doubles the energy (Fig. 2(c)).

Energy Consumption. Fig. 2(c) shows how a DTLS handshake consumes more energy with a larger wakeup interval (longer sleep periods): increasing `CI` requires a transmitter to send a longer preamble. At PDR=100%, this increase is linear. To put this energy into perspective, a DTLS handshake cost of 29.05 mJ (`CI` =500 ms, PDR=100%) represents a consumption 5 orders of magnitude lower than the energy available in a pair of AA batteries. A single DTLS handshake has hence a negligible effect on the constrained node lifetime. The cost of completing a single DTLS handshake might be more prohibitive for energy harvesting nodes with small rechargeable batteries, For example, ST GreenNet [15] nodes have a 20 mAh battery, or 201.6 J at 2.8 V. In this context, a single 29.05 mJ handshake accounts for 0.0144% of the maximum available energy.

Packet overhead. Once the DTLS session has been established, DTLS with AES_CCM_8 cipher adds 29 bytes to each datagram (including an 8-byte nonce and 8-byte authentication tag), which represents 22.8% of the available link-layer payload space (127 bytes). For the details on byte overhead and possible optimizations, see Raza *et al.* [7].

B. Beacon-Enabled IEEE 802.15.4 Networks

Nodes in a beacon-enabled IEEE 802.15.4 network are organized as a cluster tree: some nodes are cluster heads (or coordinators), others are leaf nodes. Cluster heads periodically send beacon frames. A beacon frame indicates the start of a Contention Active Period (CAP), during which leaf nodes associated with the emitting cluster head communicate using CSMA/CA medium access. After the CAP, and before the next beacon, all nodes switch their radio off. Because beacons are sent periodically, leaf nodes know when to wake up for the next CAP. The beacon interval (`BI`) and the duration of the CAP (`CAP`) are tunable, allowing a trade-off between the amount of data that can be exchanged, and the radio duty

[2]Packet Delivery Ratio (PDR) of a link is the percentage of frames successfully received at the receiver node. We use the same PDR for all links in the emulated network and control its level through the UDG radio model of Cooja.

cycle.

In our experiments, the DTLS client runs on a leaf node. In the single-hop case, the DTLS server runs on the cluster tree root, otherwise as a leaf node associated with the cluster tree root. We force a desired topology by tuning the parameters that specify the maximum number of associated cluster heads/leaf nodes such that the association requests are rejected until we obtain the desired topology. That is, we chain intermediate cluster head nodes between the DTLS client and the cluster tree root to obtain a linear network from 2 to 5 nodes, depending on the required number of hops. These cluster heads do not generate any application traffic; they only transmit periodic beacons and forward packets exchanged between DTLS server, associated with the cluster tree root, and the client.

We implement this protocol on a ST GreenNet node [15], an energy harvesting prototype from STMicroelectronics (ST) based on an ultra low power 32-bit ARM Cortex-M3 micro-controller (STM32L) with 32 kB of RAM and an IEEE 802.15.4-compliant radio transceiver. On top of the IEEE 802.15.4 beacon-enabled mode, the ST GreenNet nodes run an enhanced Contiki IPv6/6LoWPAN stack optimized for ultra low power consumption, suitable for energy harvesting nodes [15]. We port the tinyDTLS implementation to this stack and modify it to use an AES hardware acceleration block available within the micro-controller. We run tinyDTLS with exactly the same configuration as in Section III-A and estimate the energy consumption using Energest that measures the time spent by different components of the platform in a given state (for instance, the time CPU spent in active or low power mode; radio transceiver in RX or TX mode). We then derive energy by multiplying the values by the supply voltage and the current draw values from appropriate data sheets.

We obtain all the results in this section from measurements on real hardware. They are averaged over at least 500 handshakes and presented with a 95% confidence interval.

Overall results. Fig. 3 shows the measured average handshake duration and the energy consumption when DTLS runs on an IEEE 802.15.4 beacon-enabled network, in the single/multi-hop case, and for different values of `BI` and `CAP`.

Impact of `BI`. Results in the single-hop case (Fig. 3(a)) show how a short `BI` shortens the handshake as nodes get more frequent opportunities to transmit. Similarly, a large `CAP` gives nodes a long period to communicate; largest evaluated CAP of 122.9 ms yields shortest duration of the handshake. A larger `CAP` increases the throughput between two nodes, which means that a DTLS endpoint (client or server) can send its messages within the same CAP.

Impact of the number of hops. Fig. 3(b) presents the measured DTLS handshake duration when DTLS client and server are separated by multiple hops. As expected, the DTLS handshake duration increases almost linearly with the number of hops. For values of `BI` above 250 ms, the successful completion of a DTLS handshake between client and server multiple hops away requires the configuration of a large retransmission timeout value, even when there are no packets lost in the network. We observe handshake durations from 1.88 s to 16.6 s.

Energy Consumption. Fig. 3(c) shows the energy consumed by an ST GreenNet board running as a DTLS client during a DTLS handshake. The energy consumption only very slightly increases with `BI`, as the energy consumption of a transmitting node in beacon-enabled mode is not a function of the wakeup interval. Why the energy increases at all with `BI` is a consequence of the energy spent by the nodes when sitting idle.

C. IEEE 802.15.4e TSCH Networks

External interference and multi-path fading severely degrade the quality of a wireless link, both in indoor and outdoor deployments. The IEEE 802.15.4e-2012 standard [10] defines the Time Slotted Channel Hopping (TSCH) mode, which uses "channel hopping" to combat external interference and multi-path fading.

In a TSCH network, time is cut into timeslots, each long enough for a transmitter to send a data packet to a receiver, and for the receiver to send back an acknowledgment (ACK) indicating successful reception. L successive timeslots form a "slotframe", which continuously repeats over time. A communication schedule indicates to each node, for each slot in the slotframe, what to do (transmit, receive or sleep) and on which channel. The communication schedule can be built in a centralized or distributed fashion. The scheduler (either a centralized computer or a distributed protocol) builds and maintains the schedule in order to match the link-layer resources (timeslots scheduled between neighbor nodes) to the applications needs (number of packets per second, latency requirements). We assume that the scheduler schedules a cell (a [timeslot, channel] pair) to only a single pair of nodes, thereby avoiding self-interference in the network.

We derive the expected latency in a TSCH network analytically, and apply it to DTLS. Let C denote the number of cells scheduled in a slotframe between two nodes in a TSCH network. We consider that cells are distributed in the TSCH schedule in a uniform fashion, i.e. the probability for a cell to be assigned to the appropriate timeslot is $1/L$. Consider a single-hop communication between two nodes; we are interested in finding the average latency D that includes the time a frame spent in a node queue before its transmission and reception at the destination node.

Let random variable T denote the instant in a slotframe when a frame has been selected from a node queue for transmission. We consider T to be uniformly distributed over the slotframe length. Note that the instant T corresponds to a frame that is either self-originated, or received from a neighbor and to be forwarded. Let $X_0, X_1, \ldots, X_{C-1}$ denote random variables that correspond to the interval from instant T until the beginning of the corresponding cell slot. The average latency until the beginning of the frame transmission is the expectancy of the random variable $Y = min(X_0, X_1, ..., X_{C-1})$. Since the slotframe repeats in time, variables $X_0, X_1, \ldots, X_{C-1}$ are also uniformly distributed on

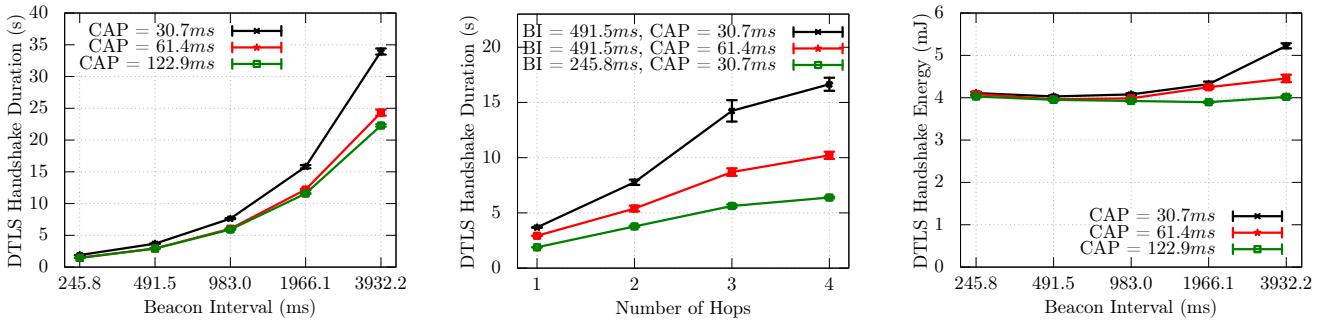

(a) DTLS handshake duration, single-hop. (b) DTLS handshake duration, multi-hop. (c) Energy, single-hop.

Fig. 3. Cost of a DTLS handshake in a beacon-enabled IEEE 802.15.4 network.

TABLE I
SINGLE-HOP DTLS HANDSHAKE DURATION IN A TSCH NETWORK.

	C = 1	**C = 2**	**C = 3**
L = 101	5.15s	3.467s	2.625s
L = 1001	50.15s	33.467s	25.125s

TABLE II
MULTI-HOP DTLS HANDSHAKE DURATION IN A TSCH NETWORK (C=1).

	H = 2	**H = 3**	**H = 4**
L = 101	10.3s	15.45s	20.6s
L = 1001	100.3s	150.45s	200.6s

$(0, L-1)$. Assuming $L >> C$, the average latency until the beginning of the frame transmission is $L/(C+1)$ timeslots. Eq. (1) expresses the single-hop latency (in timeslots), taking into account the duration of the timeslot during which the frame is transmitted, and the Packet Delivery Ratio P over the link.

$$D_{singlehop} = (1 + \frac{L}{C+1}) \cdot \frac{1}{P} \quad (1)$$

To extend Eq. (1) to the multi-hops case, we take into account the varying number of cells on each link. Considering a centralized schedule, the total latency over H intermediate hops is the sum of individual hop latencies:

$$D_{multihop} = \sum_{i=1}^{H} (1 + \frac{L}{C_i+1}) \cdot \frac{1}{P_i} \quad (2)$$

We use Eq. (1) and Eq. (2) to estimate the average duration of a DTLS handshake for typical TSCH values. We consider the same scenario as experimentally evaluated in Sections III-A and III-B (10 link-layer frames carrying DTLS messages) and a default timeslot duration of $10ms$. Tables I and II present the estimated DTLS handshake duration for typical values of L and C and ideal Packet Delivery Ratio.

We compared analytical results of TSCH with experimental results of *beacon-enabled* mode in order to find scenarios where they perform similarly. For a slotframe length of 101 timeslots, estimated handshake durations with 1, 2, and 3 dedicated cells in TSCH case, roughly correspond to *beacon-enabled* [BI = 983 ms, CAP = 61.4 ms], [BI = 491.5 ms, CAP = 30.7 ms], [BI = 491.5 ms, CAP = 61.4 ms] configurations, respectively.

IV. THE IMPACT OF MEMORY CONSTRAINTS

We have so far focused on the communication aspects of DTLS in duty-cycled networks. As the role of a DTLS server is often assumed by a constrained device, this section focuses on the effect of memory limitations on DTLS session management. RFC 7228 [16] defines three classes of constrained devices: Class 0 ($\ll$ 10 kB RAM, $\ll$ 100 kB flash), Class 1 ($\sim$ 10 kB RAM, $\sim$ 100 kB flash), Class 2 ($\sim$ 50 kB RAM, $\sim$ 250 kB flash). According to this classification, WiSMote platform is a Class 1 device, while ST GreenNet used in Section III-B falls in-between Classes 1 and 2.

Because of the way memory is allocated with embedded processors, a typical implementation statically allocates a number of DTLS "session slots", limiting the number of sessions simultaneously open. The memory footprint for a single DTLS session depends on the cipher suite and key length. The session state includes the IPv6 address and a port number of the communicating peer, its role (i.e. client or server), DTLS engine state, master secret, derived keys and other implementation specific variables. As an example, the tinyDTLS implementation uses $\sim$ 400 B of RAM per pre-shared key session, depending on the data type sizes used by different compilers, memory alignment and hardware architecture. That said, the operating system and the networking stack account for the most of the available memory. For instance, with 16 kB of RAM available on WiSMote nodes, we could only fit 3 DTLS session slots together with the full ContikiOS and the IPv6 networking stack including tinyDTLS, and a simple application for evaluation purposes.

We therefore want to determine the probability P_B that a DTLS client attempts to establish a session with a DTLS server where all DTLS session slots are already occupied. We call R the number of DTLS session slots available at a server. Let N denote the number of clients interested in establishing a session with the given server. We model the individual client rate with the exponential distribution of parameter λ, and the duration of each established session with parameter μ. Generated traffic in Erlang by each client is then $\rho = \lambda/\mu$. Under these assumptions, the blocking probability P_B of a DTLS server is simply the blocking probability of a $M/M/R/R/N$ queue, i.e. the Engset loss formula.

For instance, if we consider $R = 3$ (the case observed with WiSMote), $N = 5$ DTLS clients with $\rho = 0.5$, the blocking probability of a request is $\sim 17\%$. A DTLS server may implement different strategies for handling such requests. It may discard them or decide to close one of the open sessions in order to accommodate the newly arriving one. In the latter case, performance depends on the session closure policy, i.e. the appropriate "cache" replacement algorithm.

V. Related Work

Performance and applicability of DTLS to constrained environments has been a controversial issue for the IoT research and standardization communities. Many authors studied DTLS in this context and proposed different optimization techniques [6], [7], [17], [4], [18]. Granjal *et al.* [6] performed a comparative study on memory footprints, computational time, and the required energy between the IPsec protocol and DTLS using different cryptographic suites. They showed similar performance of IPsec and DTLS, except when DTLS is additionally used to exchange keys with the Elliptic Curve Diffie-Hellman exchange.

Hummen *et al.* [17], [4] proposed different techniques to lower the impact of the DTLS handshake on constrained devices, certificate pre-validation at the network gateway and handshake delegation to the "delegation server". Raza *et al.* [7] focused on reducing the per-datagram overhead and proposed a 6LoWPAN DTLS compression scheme. In our previous work, we studied the benefits of a stateless security architecture called OSCAR [5], based on DTLS, in order to support group communication and caching.

The recent work of Capossele *et al.* [18] explores the idea of abstracting the DTLS handshake as a CoAP resource and implementing the handshake procedure using CoAP methods. The advantage of this approach is that DTLS can leverage the reliability of confirmable CoAP messages, as well as the blockwise transfer for large messages. The drawback, however, is lost backward compatibility with the existing Internet infrastructure.

Finally, Kumar *et al.* [19] summarized DTLS memory requirements, high level communication overhead (in terms of number of messages), code size for different DTLS and cryptographic functions.

VI. Conclusion

The results of our thorough evaluation reported in this paper demonstrate surprisingly poor performance of DTLS in radio duty-cycled networks. Experiments with preamble sampling protocols show that the total duration of a DTLS handshake can be more than 50 s, depending on the Check Interval. In the case of the IEEE 802.15.4 beacon-enabled mode, we measure durations up to 35 s with the largest used Beacon Interval of ~ 4 s. Handshake duration increases linearly with the number of hops for both preamble-sampling and IEEE 802.15.4 beacon-enabled networks. We also derived the analytic expression for latency in TSCH networks and applied it to estimate the duration of DTLS handshake. For instance, in a typical TSCH network with 101 timeslots per slotframe, where DTLS server and client are radio neighbors and have a single dedicated cell for communication, it takes 5.5 s on the average to complete the handshake. This value decreases to 2.6 s for 3 dedicated cells.

Our results show that using DTLS is acceptable for applications for which a DTLS handshake is performed a limited number of times during the constrained node lifetime. For scenarios that require multiple DTLS clients per DTLS server with constrained resources, we study the blocking probability and show that it corresponds to the Engset loss formula. Applications expecting a large number of clients per DTLS server should cautiously weight the benefit of its use, against security solutions at the Network or Application layer.

References

[1] T. Dierks and E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.2*, IETF Std. RFC5246, August 2008.

[2] E. Rescorla and N. Modadugu, *Datagram Transport Layer Security Version 1.2*, IETF Std. RFC5246, January 2012.

[3] Z. Shelby, K. Hartke, and C. Bormann, *Constrained Application Protocol (CoAP)*, IETF Std. RFC7252, June 2014.

[4] R. Hummen, H. Shafagh, S. Raza, T. Voigt, and K. Wehrle, "Delegation-based Authentication and Authorization for the IP-based Internet of Things," in *Proc. of SECON*. IEEE, 2014.

[5] M. Vučinić, B. Tourancheau, F. Rousseau, A. Duda, L. Damon, and R. Guizzetti, "OSCAR: Object security architecture for the Internet of Things," *Elsevier Ad Hoc Networks*, vol. 32, pp. 3–16, 2015.

[6] J. Granjal, E. Monteiro, and J. S. Silva, "On the Effectiveness of End-to-End Security for Internet-Integrated Sensing Applications," in *Proc. GreenCom*. IEEE, 20-23 November 2012, pp. 87–93.

[7] S. Raza, H. Shafagh, K. Hewage, R. Hummen, and T. Voigt, "Lithe: Lightweight Secure CoAP for the Internet of Things," *IEEE Sensors Journal*, vol. 13, no. 10, pp. 3711–3720, 2013.

[8] M. Buettner, G. V. Yee, E. Anderson, and R. Han, "X-MAC: A Short Preamble MAC Protocol for Duty-cycled Wireless Sensor Networks," in *Proc. SenSys*. New York, NY, USA: ACM, 2006, pp. 307–320.

[9] *802.15.4-2011: IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) 1: MAC sublayer*, IEEE Std., 2011.

[10] *802.15.4e-2012: IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer*, IEEE Std., 2012.

[11] L. Seitz, S. Gerdes, G. Selander, M. Mani, and S. S. Kumar, *ACE use cases*, IETF Std. draft-seitz-ace-usecases-03 [work-in-progress], March 2015.

[12] A. Dunkels, "The ContikiMAC Radio Duty Cycling Protocol," Swedish Institute of Computer Science (SICS), Tech. Rep., 2011.

[13] "The Contiki OS," http://www.contiki-os.org/.

[14] F. Osterlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt, "Cross-Level Sensor Network Simulation with COOJA," in *Proceedings of IEEE LCN*, 2006.

[15] L.-O. Varga, G. Romaniello, M. Vučinić, M. Favre, A. Banciu, R. Guizzetti, C. Planat, P. Urard, M. Heusse, F. Rousseau, O. Alphand, E. Duble, and A. Duda, "GreenNet: an Energy Harvesting IP-Enabled Wireless Sensor Network," *IEEE Internet of Things Journal*, 2015.

[16] C. Bormann, M. Ersue, and A. Keranen, *Terminology for Constrained-Node Networks*, IETF Std. RFC7228, May 2014.

[17] R. Hummen, J. H. Ziegeldorf, H. Shafagh, S. Raza, and K. Wehrle, "Towards Viable Certificate-based Authentication for the Internet of Things," in *Proc. HotWiSec Workshop*. New York, NY, USA: ACM, 2013, pp. 37–42.

[18] A. Capossele, V. Cervo, G. De Cicco, and C. Petrioli, "Security as a CoAP Resource: an Optimized DTLS Implementation for the IoT," in *Proc. ICC*. London, UK: IEEE, 6-8 June 2015.

[19] S. S. Kumar, S. L. Keoh, and H. Tschofenig, *A Hitchhiker's Guide to the (Datagram) Transport Layer Security Protocol for Smart Objects and Constrained Node Networks*, IETF Std. draft-ietf-lwig-tls-minimal-01 [work-in-progress], March 2014.